

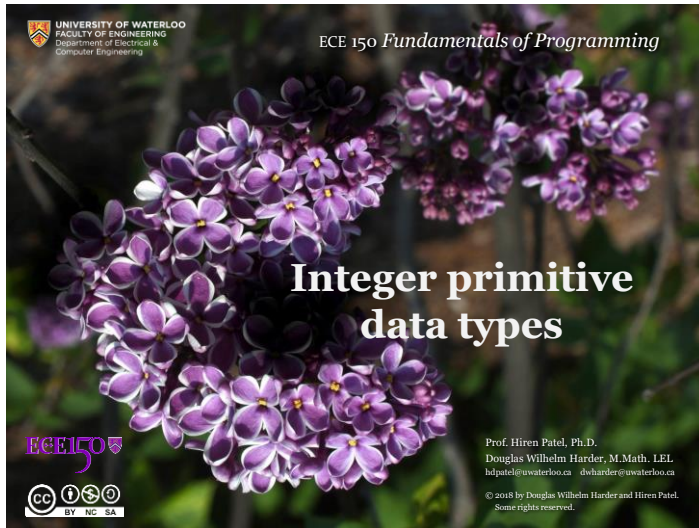
UNIVERSITY OF WATERLOO
FACULTY OF ENGINEERING
Department of Electrical & Computer Engineering

ECE 150 *Fundamentals of Programming*

Integer primitive data types

Prof. Hiren Patel, Ph.D.
Douglas Wilhelm Harder, M.Math. LEL
hdpatel@uwaterloo.ca dwharder@uwaterloo.ca

© 2018 by Douglas Wilhelm Harder and Hiren Patel.
Some rights reserved.



ECE150

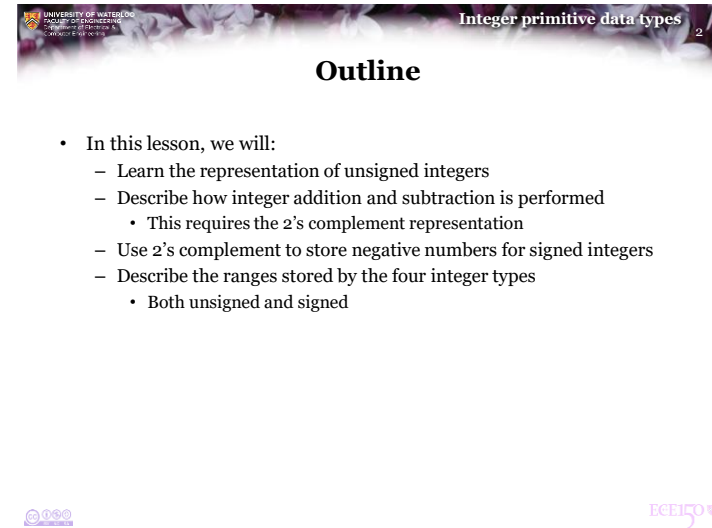
CC BY NC SA

UNIVERSITY OF WATERLOO
FACULTY OF ENGINEERING
Department of Electrical & Computer Engineering

Integer primitive data types

Outline

- In this lesson, we will:
 - Learn the representation of unsigned integers
 - Describe how integer addition and subtraction is performed
 - This requires the 2's complement representation
 - Use 2's complement to store negative numbers for signed integers
 - Describe the ranges stored by the four integer types
 - Both unsigned and signed



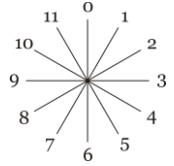
ECE150

UNIVERSITY OF WATERLOO
FACULTY OF ENGINEERING
Department of Electrical & Computer Engineering

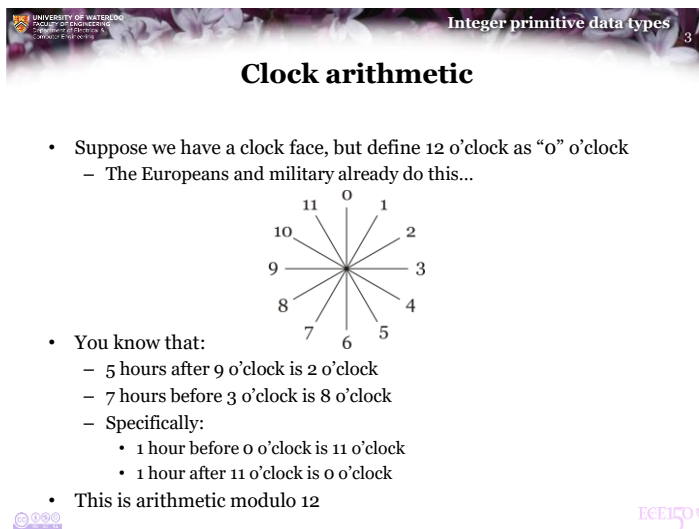
Integer primitive data types

Clock arithmetic

- Suppose we have a clock face, but define 12 o'clock as "0" o'clock
 - The Europeans and military already do this...



- You know that:
 - 5 hours after 9 o'clock is 2 o'clock
 - 7 hours before 3 o'clock is 8 o'clock
 - Specifically:
 - 1 hour before 0 o'clock is 11 o'clock
 - 1 hour after 11 o'clock is 0 o'clock
- This is arithmetic modulo 12



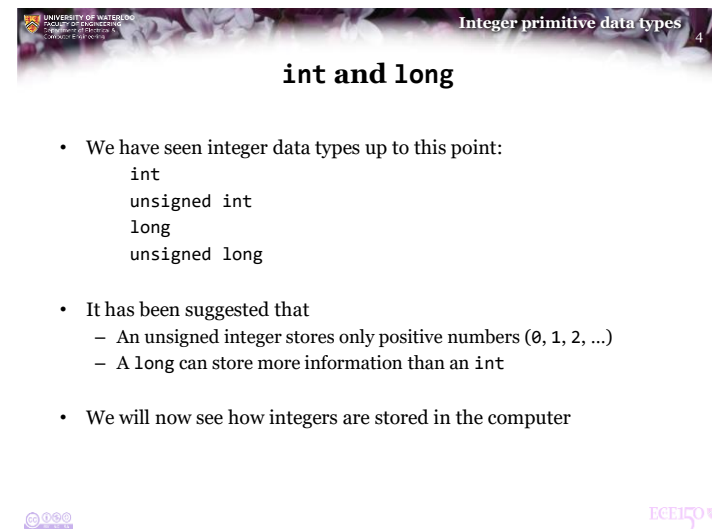
ECE150

UNIVERSITY OF WATERLOO
FACULTY OF ENGINEERING
Department of Electrical & Computer Engineering

Integer primitive data types

int and long

- We have seen integer data types up to this point:
 - int
 - unsigned int
 - long
 - unsigned long
- It has been suggested that
 - An unsigned integer stores only positive numbers (0, 1, 2, ...)
 - A long can store more information than an int
- We will now see how integers are stored in the computer



ECE150

UNIVERSITY OF WATERLOO
FACULTY OF ENGINEERING
FACULTY OF INFORMATION SYSTEMS
Integer primitive data types 5

Binary representations

- We have already described binary numbers
 - On the computer, all integers are stored in binary
 - Thus, to store each of these numbers, we must store the corresponding binary digits (bits):

3	11	2
42	101010	6
616	1001101000	10
299792458	10001110111100111100001001010	29

- To store a googol (10^{100}), we must store 333 bits:
 10010010010011010110100100101100101001100001101111100111...
 01011000010110010011110000100110001001100111000001011111...
 1001110001010110011100100000100011100010000100011010011...
 111001010101011001001000011000010001010100000101110100...
 011110001000...
 000



UNIVERSITY OF WATERLOO
FACULTY OF ENGINEERING
FACULTY OF INFORMATION SYSTEMS
Integer primitive data types 6

Storage

- Do we store as many bits as are necessary?
 - You could, but this would be exceedingly difficult to manage
- Instead, each primitive data type has a fixed amount of storage
 - 8 bits are defined as 1 byte
 - All data types are an integral number of bytes
 - Usually 1, 2, 4, 8 or 16 bytes
 - Because we use binary, powers of 2 are very common:

Exponent	Decimal	Binary
2^0	1	1
2^1	2	10
2^2	4	100
2^3	8	1000
2^4	16	10000
2^5	32	100000
2^6	64	1000000



UNIVERSITY OF WATERLOO
FACULTY OF ENGINEERING
FACULTY OF INFORMATION SYSTEMS
Integer primitive data types 7

unsigned int

- A variable is declared `unsigned int` is allocated four bytes
 - 4 bytes is $4 \times 8 = 32$ bits
 - 32 different 1s and 0s can be stored
 - The smallest and largest:
 00000000000000000000000000000000
 11111111111111111111111111111111
 - The smallest represents 0
 - The largest is one less than
 10000000000000000000000000000000
32 zeros
 - This equals 2^{32} , thus, the largest value that can be stored as an `unsigned int` is $2^{32} - 1 = 4294967295$
 - Approximately 4 billion



UNIVERSITY OF WATERLOO
FACULTY OF ENGINEERING
FACULTY OF INFORMATION SYSTEMS
Integer primitive data types 8

unsigned short

- Sometimes, you don't need to store numbers this large
- Variables declared `unsigned short` are allocated two bytes
 - 2 bytes is $2 \times 8 = 16$ bits
 - 16 different 1s and 0s can be stored
 - The smallest and largest:
 0000000000000000
 1111111111111111
 - The smallest represents 0
 - The largest is one less than
 1000000000000000
16 zeros
 - This equals 2^{16} , thus, the largest value that can be stored as an `unsigned int` is $2^{16} - 1 = 65535$



UNIVERSITY OF WATERLOO
FACULTY OF ENGINEERING
FACULTY OF INFORMATION TECHNOLOGY

Integer primitive data types

13

Example

- Generally, however, we display the bytes in memory as a column of bytes, the values of which are concatenated

```
#include <iostream>

// Function declarations
int main();

// Function definitions
int main() {
    unsigned short a(42);
    unsigned int b(207500);
    unsigned long c(299792458);

    std::cout << (a + b + c) << std::endl;

    return 0;
}
```

8 bytes for c

```
00000000
00000000
00000000
00000000
00010001
11011110
01111000
01001010
00000000
00000011
00101010
10001100
00000000
00101010
```

4 bytes for b

2 bytes for a



UNIVERSITY OF WATERLOO
FACULTY OF ENGINEERING
FACULTY OF INFORMATION TECHNOLOGY

Integer primitive data types

14

Wasted space?

- If an integer does not use all the bytes, the remaining bits are nevertheless allocated until the variable goes out of scope
 - In general-purpose computing, this is often not a problem
 - This is a critical issue, however, in embedded systems
 - More memory:
 - Costs more
 - Uses more power
 - Produces more heat



UNIVERSITY OF WATERLOO
FACULTY OF ENGINEERING
FACULTY OF INFORMATION TECHNOLOGY

Integer primitive data types

15

Determining the size of a type

- We have said short, int and long are 2, 4 and 8 bytes
 - This is true on most every general-purpose computer
- Unfortunately, the C++ specification doesn't require this
 - Fortunately, the sizeof operator gives you this information

```
#include <iostream>
int main();
int main() {
    std::cout << "An 'unsigned short' occupies "
              << sizeof ( unsigned short ) << " bytes" << std::endl;
    std::cout << "An 'unsigned int' occupies "
              << sizeof ( unsigned int ) << " bytes" << std::endl;
    std::cout << "An 'unsigned long' occupies "
              << sizeof ( unsigned long ) << " bytes" << std::endl;
    return 0;
}
```

Output on *eclinux*:

```
An 'unsigned short' occupies 2 bytes
An 'unsigned int' occupies 4 bytes
An 'unsigned long' occupies 8 bytes
```



UNIVERSITY OF WATERLOO
FACULTY OF ENGINEERING
FACULTY OF INFORMATION TECHNOLOGY

Integer primitive data types

16

Memory and initial values

- Question:
 - What happens if the initial value cannot be stored?

```
#include <iostream>

int main();

int main() {
    unsigned short c(299792458);
    std::cout << "The speed of light is " << c
              << " m/s." << std::endl;

    return 0;
}
```



UNIVERSITY OF WATERLOO
FACULTY OF ENGINEERING
FACULTY OF INFORMATION AND COMMUNICATIONS

Integer primitive data types 17

Memory and initial values

- Fortunately, you get a warning:

```
example.cpp: In function 'int main()':
example.cpp:6:31: warning: narrowing conversion of â299792458â from 'int' to
'short unsigned int' inside { } [-Wnarrowing]
    unsigned short c(299792458);
                          ^
example.cpp:6:31: warning: large integer implicitly truncated to unsigned
type [-Woverflow]
```
- It still compiles and executes:

The speed of light is 30794 m/s.



UNIVERSITY OF WATERLOO
FACULTY OF ENGINEERING
FACULTY OF INFORMATION AND COMMUNICATIONS

Integer primitive data types 18

Memory and initial values

- Where does 30794 come from?

c requires 29 bits

```
00010001110111100111100001001010
```

Only 16 bits are allocated
- The binary number `0b111100001001010` equals 30794 in base 10



UNIVERSITY OF WATERLOO
FACULTY OF ENGINEERING
FACULTY OF INFORMATION AND COMMUNICATIONS

Integer primitive data types 19

Memory and initial values

- Important:

*All unsigned integers are stored:
modulo 2^{16} for unsigned short
modulo 2^{32} for unsigned int
modulo 2^{64} for unsigned long*



UNIVERSITY OF WATERLOO
FACULTY OF ENGINEERING
FACULTY OF INFORMATION AND COMMUNICATIONS

Integer primitive data types 20

Memory and arithmetic

- What happens if the sum, difference or product of two integers exceeds what can be stored?

```
#include <iostream>

int main();

int main() {
    unsigned short m1(40000), m2(42000);
    int n1(40000), n2(42000);
    unsigned short sum(m1 + m2), diff(m1 - m2), prod(m1*m2);

    std::cout << sum << "\t" << (n1 + n2) << std::endl;
    std::cout << diff << "\t" << (n1 - n2) << std::endl;
    std::cout << prod << "\t" << (n1*n2) << std::endl;

    return 0;
}
```

Output:

16464	82000
63536	-2000
50176	168000000



UNIVERSITY OF WATERLOO
FACULTY OF ENGINEERING
FACULTY OF DESIGN
FACULTY OF MATHEMATICS

Integer primitive data types 21

Memory and arithmetic

- Let's look at the actual values and the evaluated results:

16464	010000001010000
82000	1010000001010000
63536	111110000110000
-2000	-0000011111010000
50176	110001000000000
168000000	110010000100010110001000000000

- For the sum and product, the result ignores the higher-order bits
 - The negative number is a little odd....



UNIVERSITY OF WATERLOO
FACULTY OF ENGINEERING
FACULTY OF DESIGN
FACULTY OF MATHEMATICS

Integer primitive data types 22

Memory and arithmetic

- What happens if the sum, difference or product of two integers exceeds what can be stored?

```
#include <iostream>
int main();
int main() {
    unsigned short smallest{0}, largest{65535};

    std::cout << "Smallest: " << smallest << std::endl;
    std::cout << "Largest: " << largest << std::endl;
    --smallest;
    ++largest;
    std::cout << "Smallest minus 1: " << smallest << std::endl;
    std::cout << "Largest plus 1: " << largest << std::endl;

    return 0;
}
```

Output:
Smallest: 0
Largest: 65535
Smallest minus 1: 65535
Largest plus 1: 0




UNIVERSITY OF WATERLOO
FACULTY OF ENGINEERING
FACULTY OF DESIGN
FACULTY OF MATHEMATICS

Integer primitive data types 23

Memory and arithmetic

- Important:
 - All unsigned integers arithmetic is performed: modulo 2¹⁶ for unsigned short modulo 2³² for unsigned int modulo 2⁶⁴ for unsigned long*

- This is similar to all clock arithmetic being performed modulo 12




UNIVERSITY OF WATERLOO
FACULTY OF ENGINEERING
FACULTY OF DESIGN
FACULTY OF MATHEMATICS

Integer primitive data types 24

Addition

- Addition is easy:
 - Like in elementary school, line them up and occasionally you require a carry in the next column:
 - The rules are:
 - 0 + 0 → 0
 - 0 + 1 → 1
 - 1 + 1 → 10 → 0 with a carry of 1
 - 1 + 1 + 1 → 11 → 1 with a carry of 1
 - For example, adding two unsigned short:



Addition

- What if we go over? Adding these two unsigned short:

$$\begin{array}{r}
 1 \ 1 \ 1 \ 1 \ 1 \ 1 \\
 1101000111010110 \quad \leftarrow 53718 \\
 + 1001101011000100 \quad \leftarrow 39620 \\
 \hline
 10110110010011010 \quad \leftarrow 93338
 \end{array}$$

- The additional bit is discarded—addition is calculated modulo 2^{16}
 - Thus, the answer is 110110010011010 which is 27802



Subtraction

- Subtraction is more difficult:
 - Like in elementary school, you learned to “borrow”, but borrowing may require you to look way ahead:


$$\begin{array}{r}
 0100000001010000 \\
 - 0001101011000101 \\
 \hline
 ?
 \end{array}$$

- Our salvation: we are performing arithmetic modulo 65536



Subtraction

- Going back to the clock:
 - Subtracting 10 is the same as adding 2
 - Subtracting 4 is the same as adding 8
 - Subtracting 9 is the same as adding 3
- Thus, to subtract n , add $12 - n$



- In our case, to subtract n , add $65536 - n$

$$\begin{array}{r}
 0100000001010000 \quad \leftarrow 16464 \\
 - 0001101011000101 \\
 \hline
 ? \\
 + 1110010100111011 \quad \leftarrow 58683 \\
 + 10010010110001011
 \end{array}$$

- The answer is 0010010110001011 \leftarrow 9611



Subtraction

- The million-dollar question:
 - How do you calculate $65536 - n$???
- Subtract any number from 999999999999, no borrows are needed

$$\begin{array}{r}
 999999999999 \\
 - 5501496383498 \\
 \hline
 4498503616501
 \end{array}$$

- Thus, to calculate $1000000000000 - n$, instead calculate $(1000000000000 - 1) - n + 1 = (999999999999 - n) + 1$

- For example:

$$\begin{array}{r}
 1000000000000 \\
 - 5501496383498 \\
 \hline
 4498503616502
 \end{array}$$

This is called the base-10 complement or “10’s complement”
 – this is how older adding machines performed subtraction





Subtraction

- In binary, the equivalent is base-2 complement or “2’s complement”
 - To calculate $65536 - 1970$, calculate $(65536 - 1970) + 1$:

$$\begin{array}{r}
 1111111111111111 \\
 - 0000011110110010 \\
 1111100001001101 \\
 + \underline{1} \\
 1111100001001110
 \end{array}$$

- Thus, to calculate $2018 - 1970$, just add the 2’s complement of 1970 to 2018:

$$\begin{array}{r}
 0000011111100010 \\
 + 1111100001001110 \\
 \hline
 1000000000110000
 \end{array}$$

- This is the binary representation of $48 = 2^5 + 2^4 = 32 + 16$
 - Remember, we ignore the leading 1



2’s complement

- To calculate the 2’s complement:
 - Complement all of the bits in the number
 - This includes leading zeros
 - Add 1
- For example, the 2’s complement of the speed of light is stored as an unsigned int is

$$\begin{array}{r}
 00010001110111100111100001001010 \\
 11101110001000011000011110110101 \\
 + \underline{1} \\
 11101110001000011000011110110110
 \end{array}$$



2’s complement

- There is a faster way to compute it without the addition:
 - Scan from right-to-left
 - Find the first 1, and then flip each bit to the left of that
- The 2’s complement of each of the following is given below it

$$\begin{array}{r}
 1011011111011111 \\
 0100100000100001
 \end{array}$$

$$\begin{array}{r}
 1010111111100000 \\
 0101000000100000
 \end{array}$$

$$\begin{array}{r}
 0000100100101100 \\
 1111011011010100
 \end{array}$$



2’s complement

- The 2’s complement of 0 stored as an unsigned int is
- $$\begin{array}{r}
 00000000000000000000000000000000 \\
 11111111111111111111111111111111 \\
 + \underline{1} \\
 10000000000000000000000000000000
 \end{array}$$
- This makes sense: any number minus zero is unchanged



2's complement

- The 2's complement algorithm is self-inverting:
 - If n is a number, then $2^{16} - (2^{16} - n) = n$
 - The 2's complement of the 2's complement of a number is the number itself

```

1110110010111110
0001001101000001
+ ----- 1
0001001101000010
1110110010111101
+ ----- 1
1110110010111110
    
```

- That is, $f^{-1} = f$ or $f(f(n)) = n$



Memory and arithmetic

- Try it yourself:

```

#include <stdio.h>
int main() {
    int x = 10;
    int y = 20;
    int z = x + y;
    printf("x + y = %d\n", z);
    return 0;
}
    
```



Summary so far

- We have the following:
 - Unsigned integers are stored as either 1, 2, 4 or 8 bytes
 - The value is stored in the binary representation

Type	Bytes	Bits	Range	Approximate Range
unsigned char	1	8	0, ..., $2^8 - 1$	0, ..., 255
unsigned short	2	16	0, ..., $2^{16} - 1$	0, ..., 65535
unsigned int	4	32	0, ..., $2^{32} - 1$	0, ..., 4.3 billion
unsigned long	8	64	0, ..., $2^{64} - 1$	0, ..., 18 quintillion

- You should not memorize the exact ranges



Useful tool...

- Note that $2^{10} = 1024$, so $2^{10} \approx 1000 = 10^3$
 - We can use this to estimate magnitudes:
 - $2^{12} = 2^2 2^{10} \approx 4 \times 1000 = 4000$
 - $2^{16} = 2^6 2^{10} \approx 64 \times 1000 = 64000$
 - $2^{24} = 2^4 2^{20} = 2^4 (2^{10})^2 \approx 16 \times 1000^2 = 16 \text{ million}$
 - $2^{32} = 2^2 2^{30} = 2^2 (2^{10})^3 \approx 4 \times 1000^3 = 4 \text{ billion}$
 - This approximation will underestimate by approximately 2%



Signed types

- We've seen that short, int and long all allows you to store both positive and negative integers
 - How do we store such negative numbers?
- Because we have two choices (positive or negative), we could use one bit to represent the *sign*: 0 for positive, 1 for negative
 - For example:


The sign bit

32767	0111111111111111	
2	000000000000010	
1	000000000000001	
0	000000000000000	
-0 ?	100000000000000	-0 = 0, so do we have two zeros?
-1	100000000000001	
-2	100000000000010	
-32768	1111111111111111	



Signed types

- This is similar to marking the hours of a clock as follows:




- Unfortunately, this leads to ugly arithmetic operations...
 - 1 + 1 = 0 or -0, but 7 + 1 = 8
 - 5 + 2 = -3, but 11 + 2 = 1



Signed types

- A better solution:



- Note that
 - 1 + 1 = 0, but also 11 + 1 = 0
 - 5 + 2 = -3, but also 7 + 2 = 9, which we are equating to -3



Signed integers

- Here is a workable solution:
 - If the leading bit is 0:
 - Assume the remainder of the number is the integer represented
 - For short, this includes
 - 0000000000000000 0
 - 0111111111111111 $2^{15} - 1 = 32767$
 - This includes 2^{15} different positive numbers
 - If the leading bit is 1:
 - Assume the number is negative and its magnitude can be found by applying the 2's complement algorithm
 - Recall the 2's complement algorithm is self-inverting



Signed integers

- For negative numbers stored as a short:

$$\begin{array}{r} 100000000000000 \\ 011111111111111 \\ + \frac{1}{100000000000000} \\ \hline 100000000000000 \end{array}$$

- This is the representation of the largest negative number: -2^{15}


$$\begin{array}{r} 111111111111111 \\ 000000000000000 \\ + \frac{1}{000000000000001} \\ \hline 000000000000001 \end{array}$$

- This is the representation of the smallest negative number: -1



Signed integers

- Here, you can compare these two techniques
 - In both cases, we go from $-12/2$ to $12/2 - 1$ and $-2^{16}/2$ to $2^{16}/2 - 1$



Signed integers

- For example, 111111111010110 is a negative short

$$\begin{array}{r} 111111111010110 \\ 000000000101001 \\ + \frac{1}{000000000101010} \\ \hline 000000000101010 \end{array}$$

- Thus, it represents -42
- Let's calculate $-42 + 91 = 49$ and $-42 - 91 = -133$:

$$\begin{array}{r} 111111111010110 \\ + 000000000101011 \\ \hline 10000000000110001 \end{array} \leftarrow 49$$

$$\begin{array}{r} 111111111010110 \\ + 111111111010010 \\ \hline 11111111101111011 \end{array} \leftarrow -133$$

$10000101 \leftarrow 133$



Summary

- To summarize:
 - Integer types are stored as either 1, 2, 4 or 8 bytes
 - Negative numbers are stored in the 2's complement representation

Type	Bytes	Bits	Range	Approximate Range
unsigned char	1	8	$0, \dots, 2^8 - 1$	0, ..., 255
unsigned short	2	16	$0, \dots, 2^{16} - 1$	0, ..., 65535
unsigned int	4	32	$0, \dots, 2^{32} - 1$	0, ..., 4.3 billion
unsigned long	8	64	$0, \dots, 2^{64} - 1$	0, ..., 18 quintillion
signed char	1	8	$-2^7, \dots, 2^7 - 1$	-128, ..., 127
short	2	16	$-2^{15}, \dots, 2^{15} - 1$	-32768, ..., 32767
int	4	32	$-2^{31}, \dots, 2^{31} - 1$	-2.15 billion, ..., 2.15 billion
long	8	64	$-2^{63}, \dots, 2^{63} - 1$	-9 quintillion, ..., 9 quintillion





Warning

- While common, the C++ standard does not require these sizes:
 - Each compiler may choose sizes so long as the following are true:


```
assert( sizeof( char ) == 1 );
assert( sizeof( short ) >= 2 ); // At least 16 bits
assert( sizeof( int ) >= sizeof( short ) );
// At least as large as 'short'
assert( sizeof( long ) >= 4 ); // At least 32 bits
assert( sizeof( long long ) >= 8 ); // At least 64 bits
```
- In GNU g++, the sizes are as we have described in this slide deck
- In Microsoft Visual Studio, however:
 - A long is only four bytes (same as int)
 - A long long is eight bytes
 - We do not use long long in this course
 - You may have to use it if you program in Visual Studio



Summary

- Following this lesson, you now
 - Understand the representation of unsigned integers
 - Know how to perform subtraction using 2's complement
 - Similar to 10's complement used a century ago
 - Understand that signed integers store negative numbers in their 2's complement representation
 - Know that char is actually just an integer type
 - It can be interpreted as a printable character if necessary
 - Understand the ranges stored by char, short, int and long



References

- [1] Wikipedia
[https://en.wikipedia.org/wiki/Integer_\(computer_science\)](https://en.wikipedia.org/wiki/Integer_(computer_science))
https://en.wikipedia.org/wiki/Two%27s_complement



Acknowledgments

Theresa DeCola and Charlie Liu.





Colophon

These slides were prepared using the Georgia typeface. Mathematical equations use Times New Roman, and source code is presented using Consolas.

The photographs of lilacs in bloom appearing on the title slide and accenting the top of each other slide were taken at the Royal Botanical Gardens on May 27, 2018 by Douglas Wilhelm Harder. Please see

<https://www.rbg.ca/>

for more information.



Disclaimer

These slides are provided for the ECE 150 *Fundamentals of Programming* course taught at the University of Waterloo. The material in it reflects the authors' best judgment in light of the information available to them at the time of preparation. Any reliance on these course slides by any party for any other purpose are the responsibility of such parties. The authors accept no responsibility for damages, if any, suffered by any party as a result of decisions made or actions based on these course slides for any other purpose than that for which it was intended.

